

REMARKS

Pending Claims

Claims 1, 5, 6, and 8-24 are still pending, of which claims 1, 15, and 16 are independent. The applicant is not requesting any amendment to the claims presented in the response to the previous Office action.

Claim Rejections Under 35 U.S.C. 103(a)

The Examiner has rejected all of the pending claims as being unpatentable under 35 U.S.C. 103(a) over Yates, et al. (US 6,549,959 - "*Yates*") in view of *Krishnaswamy* (US 6,308,318 - "*Krishnaswamy*").

With respect to each of the independent claims, the Examiner wrote:

Yates fails to disclose specifically delaying handling of each asynchronous sensed exception, at any point in the translated target instruction, from the point in the translated target instruction sequence at which the asynchronous exception was sensed.

Krishnaswamy discloses delaying handling of each asynchronous sensed exception (col 3, lines 45-67), at any point in the translated target instruction (col 3, lines 45-67, col 4, lines 1-25), from the point in the translated target instruction sequence at which the asynchronous exception was sensed (col 3, lines 45-55, col 4, lines 11-20, col 5, lines 10-17, col 6, lines 1-15).

Krishnaswamy's system delays handling of asynchronous exceptions that occur during execution of target instructions, but handling occurs only if execution has been interrupted at or has reached a "GAR point," which *Krishnaswamy* defines as "a location in the translated code where state can be recovered, i.e., where the application's context can be correctly reconstructed" (col. 7, lines 41-44). Each such GAR point must have been specifically "programmed into the translated code" (col. 4, lines 7-8, col. 6, lines 63-64), "normally [] only separated by approximately 10 or 15 instructions" (col. 4, lines 19-20). As Figure 4 and the accompanying text show, *Krishnaswamy's* system does not provide for

the case where no GAR point is located in the translated sequence that is being executed when an asynchronous exception occurs.

According to *Krishnaswamy*, GAR points may be programmed anywhere into a target instruction sequence. See col. 7, lines 26-39 (emphasis added):

One of the advantages of the present invention is that GAR points in the translated code do not have to be at any particular locations in the code. For example, it is not necessary for GAR points to be located at the end of each block of code, as is the case in other known dynamic translation systems. This is because the present invention does not require the use of GAR points which comprise instructions that check to determine whether an asynchronous exception has occurred. In accordance with the present invention, the GAR points may be at, for example, the beginning of a translation, the head of a loop in the interior of a translation, or before an interruptible system call. Furthermore, in accordance with the present invention, the GAR points can be arbitrarily located.

In *Krishnaswamy's* preferred embodiment, GAR points are placed at the end of each "basic block of instructions" (col. 6, lines 16-17), where a "basic block is a sequence of instructions which ends with a branch and which contains no other branches" (col. 6, lines 18-19). Note that *Krishnaswamy's* "block" is therefore not the same as a "translated target instruction sequence" as defined in the present application, in particular, in the independent claims, in which "each instruction in the source instruction sequence [is] converted into a corresponding translated target instruction sequence, which may consist of a single target instruction": *Krishnaswamy's* basic block could thus correspond to translations of more than one source instruction, whereas each translated target instruction sequence in the applicant's invention does not.

The main goal of *Krishnaswamy's* method and apparatus is to avoid having to check for the occurrence of an asynchronous exception (col. 3, lines 23-37; col. 4, lines 20-24; col. 6, lines 36-44; etc.). The main goal of the applicant's invention is to guarantee precise handling of sensed asynchronous exceptions while still allowing forward progress in the translated instruction sequence that is being executed when the exception arises (see, for example, the specification on page 25, lines 9-23). These different goals are reflected in the different mechanisms that *Krishnaswamy* and the applicant provide.

Because *Krishnaswamy's* system, when running, cannot know in advance where a GAR point may be located in a currently executing code block when an asynchronous exception occurs, his interpreter must frequently check whether execution is at a GAR point; the frequency can be as often as every instruction, but may be as often as every *branch* instruction. This is not an issue for the applicant's invention. Rather, the applicant's specification clearly defines and the independent claims clearly and specifically recite what is meant by a "source instruction sequence" and a "target instruction sequence," (which is not the same as *Krishnaswamy's* "blocks") and the claims explicitly recite, upon occurrence and sensing of an exception and determining whether each sensed exception is asynchronous, *"resuming to completion the execution of the target instruction sequence in binary translation"* (not via interpretation) and thus *"delaying handling of each asynchronous sensed exception until, and no later than, completion of execution of the target instruction sequence corresponding to the current source instruction when the asynchronous exception is sensed."* *Krishnaswamy* lacks this guarantee.

The applicant's invention thus does not need to frequently check whether execution is at a programmed-in "GAR point" since it uses the properties of binary translation of a source instruction to decide when to handle the asynchronous exception: As the independent claim recite, the applicant's invention generates a *"mapping between the beginning and ending addresses of each source instruction and its corresponding translated target instruction sequence"* so it will always know when it has *"resum[ed] to completion the execution of the target instruction sequence in binary translation"* and can begin to handle the asynchronous exception.

Krishnaswamy discloses no such mechanism. Indeed, as long as it is possible to reconstruct state before the end of a translated instruction sequence, *Krishnaswamy* could have a GAR point within the sequence and have no guarantee that the entire sequence would ever complete execution at all; for example, handling of the asynchronous exception could cause a *synchronous* fault, or some other event that precludes continued execution of the interrupted target instruction sequence. This would never destroy source instruction atomicity in the applicant's invention as claimed

whereas it could in *Krishnaswamy's*. In other words, handling of an asynchronous exception might cause a more serious failure in the applicant's system, but at least the failure would occur "naturally," that is, at a *source* instruction boundary, not at some intermediate execution point.

The independent claims therefore define the invention to include several features that are lacking in *Krishnaswamy* and that provide a clear and important performance advantage, namely, guaranteed atomicity of execution of translated source instructions when an asynchronous exception occurs, with no need for special "GAR points" to be inserted in the translation. As explained in previous Office action responses, and as the Examiner himself concedes, Yates also fails to disclose specifically delaying handling of each asynchronous sensed exception in the manner claimed by the applicant. Consequently, neither *Yates* nor *Krishnaswamy*, and thus no hypothetical combination of these two systems, has the novel features of the applicant's invention as claimed. The independent claims 1, 15, and 16 should therefore now be allowable.

The Examiner also rejected all the dependent claims in view of either *Yates* alone, or the same hypothetical combination of both *Yates* and *Krishnaswamy*. All of these dependent claims inherit the limitations of their respective base claim, that is, claim 1, 15 or 16. Since the independent base claims should now be allowable, so should the dependent claims.

Request to Withdraw Finality of Rejection

As explained above, the claims should be allowed without further amendment. Nonetheless, if the Examiner still feels there is a substantial issue preventing allowance, the applicant respectfully requests not only that the finality of the 9 September 2004 office action be withdrawn, but also that he be given another office action on the merits without the need to file yet another continuation application.

M.P.E.P 707.07(g) states:

Piecemeal examination should be avoided as much as possible. The examiner ordinarily should reject each claim on all valid grounds available,

avoiding, however, undue multiplication of references.

The previous Office action (mailed 22 April 2004) was the third that the Examiner issued. In it, the Examiner cited *Robinson* (as a secondary reference in combination with *Yates*) for the first time, even though *Robinson* was available to the Examiner in preparing the first Office action over a year ago. Wanting to avoid delay in prosecution, the applicant explained the patentability of the invention over *Yates* and *Robinson* well within the three-month shortened statutory period for response.

Now it appears that the Examiner did not consider *Robinson* to be the best reference after all, since he now rejects the claims based in part on a different reference – *Krishnaswarty* – that also was available to the Examiner in preparing the first Office action, and certainly instead of or in addition to the *Robinson* reference previously cited.

M.P.E.P. 904.03 (Conducting the Search) states (emphasis added):

It is a prerequisite to a **speedy and just determination** of the issues involved in the examination of an application that a careful and comprehensive search, commensurate with the limitations appearing in the most detailed claims in the case, be made in preparing the first action on the merits so that the second action on the merits can be made final or the application allowed with no further searching other than to update the original search ... Applicants can facilitate a complete search by including, at the time of filing, claims varying from the broadest to which they believe they are entitled to the most detailed that they would be willing to accept.

In doing a complete search, the examiner should find and cite references that, while not needed for treating the claims, would be useful for forestalling the presentation of claims to other subject matter regarded by applicant as his or her invention, by showing that this other subject matter is old or obvious.

In selecting the references to be cited, ... the examiner is not called upon to cite all references

that may be available, **but only the "best."** (37 CFR 1.104(c).) ... **The best reference should always be the one used.**

706.07(a) (Final Rejection, When Proper on Second Action) states (emphasis added):

Under present practice, second or any subsequent actions on the merits shall be final, **except where the examiner introduces a new ground of rejection that is neither necessitated by applicant's amendment of the claims** nor based on information submitted in an information disclosure statement ...

The applicant has at every point in the prosecution of this application done everything possible to facilitate a "speedy and just determination of the issues involved in the examination," but is faced with a piecemeal presentation of references and rejections. If the Examiner does not agree that the present claims are allowable, then the applicant therefore respectfully asserts that an opportunity to respond to another Office action on the merits would greatly further the Office's stated goal of a speedy and just determination of the issues involved in the examination of this application.

Conclusion

The applicant's invention as defined in the independent claims includes at least one feature that is not found or suggested in the cited prior art, and that provides a substantial technical advantage. As such, the applicant respectfully submits that the independent claims should be allowed, as well as the remaining, dependent claims that simply further narrow the definition of the invention.

Date: 28 September 2004

34825 Sultan-Startup Rd.
Sultan, WA 98294
Phone & fax: (360) 793-6687

Respectfully submitted,

Jeffrey Pearce

Jeffrey Pearce
Reg. No. 34,729
Attorney for the Applicant